

Chapter 15

Data Reduction for Synthesis Imaging

J. THOMAS ARMSTRONG

NAVAL RESEARCH LABORATORY
WASHINGTON, D.C.

15.1 Preliminaries

I will start with a collection of ideas that underlie the rest of this Chapter. The first idea is that in optical interferometry the phase reference is the star itself, because the fringe motions are so rapid (roughly a radian of fringe phase in ~ 10 msec at visual wavelengths) and so large (tens of microns over a few seconds) that using another source as the phase reference is impossible.* This leads to a second basic notion, that the targets of current optical interferometers are stars. No other objects have enough surface brightness to act as their own phase calibrators.

Now, fringe visibilities from a stellar image are high only in the inner part of the (u, v) plane. For a star whose image is a uniform disk, $V^2(\lambda) \propto [J_1(\mathbf{x})/\mathbf{x}]^2$ i.e., $V^2(\mathbf{x})$ has the radial profile of an Airy disk, with its first zero at $\mathbf{x} = 3.83$ (refer to Table 15.1 for notation). Consequently, we arrive at a third basic notion, that there is a longest baseline length on which the fringe visibility is high enough to allow us to track the fringe. The smallest V^2 that can be tracked depends on the signal-to-noise ratio, which is a function of NV^2 (cf. Equation 14.1) and thus depends on both the brightness and angular size of the star.

If, for instance, we can track fringes when V^2 is $\simeq 0.2$, the constraint on array-element separations means that the elements must form a chain in which no link is longer than

*The Palomar Testbed Interferometer (PTI) circumvents this problem by looking at two stars simultaneously, but no other interferometers currently under development will use this technique.

Table 15.1: Notation

Symbol	Meaning
B, \mathbf{B}	Baseline length; baseline vector
FASN	Fringe amplitude SNR
$J_i(\cdot)$	Bessel function of first kind of order i
LD	Limb-darkened disk
N	Photons per integration
SNR	Signal to noise ratio
UD	Uniform disk
(u, v)	Spatial frequency coordinates; $\sqrt{u^2 + v^2} = B/\lambda$
V^2	Squared fringe visibility; loosely, “visibility”
V_*^2	True visibility of a star
V_{raw}^2	Observed visibility
$V^{(3)}$	Amplitude of triple product
\mathbf{x}	$= \pi\theta\mathbf{B}/\lambda$, so $V_{*,\text{UD}}^2 = [J_1(\mathbf{x})/\mathbf{x}]^2$ has its first zero at $\mathbf{x} = 3.83$
θ	Stellar angular diameter
λ	Wavelength
ϕ_{cl}	Closure phase (= phase of triple product)
$\vec{\rho}$	Binary separation vector

$\sim 125\lambda_{\mu\text{m}}/\theta_{\text{mas}}$. For future convenience, I will designate this greatest minimum separation as B_{track} , the longest baseline on which fringes from a given star can be detected within the coherence time of the atmosphere so they can be tracked.

The final notion is that an interferometer ought to be designed with both wavelength and baseline bootstrapping in mind, for the reasons Dave Mozurkewich went into in Chapter 14. The following discussion assumes that both are available. My discussion will also draw on several of the V^2 models that Dave generated, so you should be prepared to refer frequently to Chapter 14.

15.2 Planning the Observations

In planning the observations, you want to choose the array configuration that fits the source you’re looking at, with baseline lengths appropriate to the angular scale you’re interested in. Keeping the array phased up is a major constraint in some circumstances. You also want to observe the star at enough hour angles to get sufficient (u, v) coverage, and get enough scans on calibrator stars to monitor seeing changes.

The first thing to determine when choosing a configuration is whether the angular scale of the structure that you are interested in is larger or smaller than the stellar disk.

When the structure of interest is small compared to the stellar disk, the requirements of phasing up the array are paramount. The most obvious example of small structure is stellar surface structure, but a second is limb darkening. It is when we are observing small structure that wavelength and baseline bootstrapping are vital. In order to bootstrap, we need a configuration for which every link in the chain of array elements is shorter than, but as close as possible to, B_{track} in order to maximize the longest baseline of the array.

If the structure of interest is larger than the star, we have a “source with high visibility” in Dave Mozurkewich’s terminology (Section 14.1.2). Examples of observing large structure include measuring a binary orbit and imaging circumstellar dust or $\text{H}\alpha$ emission. For these sources, we choose the configuration that best constrains the image or the model, rather than the one that maximizes the longest baseline. In general, that means choosing the most uniform (u, v) coverage, which in turn means choosing a variety of telescope separations.

Dave touched on how to deal with the various cases of observing small and large structure. In this section, I will go into more detail with four examples of the effects of differing observing goals on the choice of configuration: determining a stellar uniform-disk diameter; observing stellar surface structure, including limb darkening; determining a binary-star separation; and observing circumstellar emission.

15.2.1 Stellar Diameters

If all that we want is an equivalent uniform disk diameter, θ_{UD} , we do not need a multi-element interferometer; we need only a single baseline of length $\sim B_{\text{track}}$. Choosing a configuration reduces to choosing a baseline long enough to partially resolve the star, but still short enough to track the fringes. If enough wavelength bootstrapping is available, the bluest channels can reach the first null of V^2 and give us a very precise measurement of θ_{UD} .

To put it in more concrete terms, I will show an example with the NPOI. Consider some observations of ϵ Gem (G8 Ib; $m_V = 4.38$, $(V - R) = 0.96$, $(R - I_C) = 0.60$; $\theta_{LD} = 4.89$ mas [Pauls 1998]) made with the NPOI on 26 March 1997 with three elements of the astrometric array. Figure 15.1 shows uncalibrated visibilities from the longest baseline and one of the short baselines, as well as uncalibrated closure phases. The baseline lengths were 18.9, 22.2, and 37.5 m; about 500 photons were received per 2-ms data frame; and uncalibrated V^2 values ranged from 0.4 at the smallest (u, v) spacings (the 18.9 m baseline at $\lambda = 850$ nm), to zero at the longest, varying with time because the projected baseline length changes as the Earth rotates. The resulting fringe amplitude SNR (FASN), derived from Equation 14.1 in Chapter 14, was ~ 8 on the shorter baselines, but was as low as ~ 1.5 on the long baseline.

If we pretend that the 18.9 m baseline data shown in Figure 15.1(a) is from our one-baseline interferometer, we have plenty of FASN to track the fringes, yet the decline in visibility with wavelength (i.e., with increasing (u, v) distance) is enough to fit a diameter to the

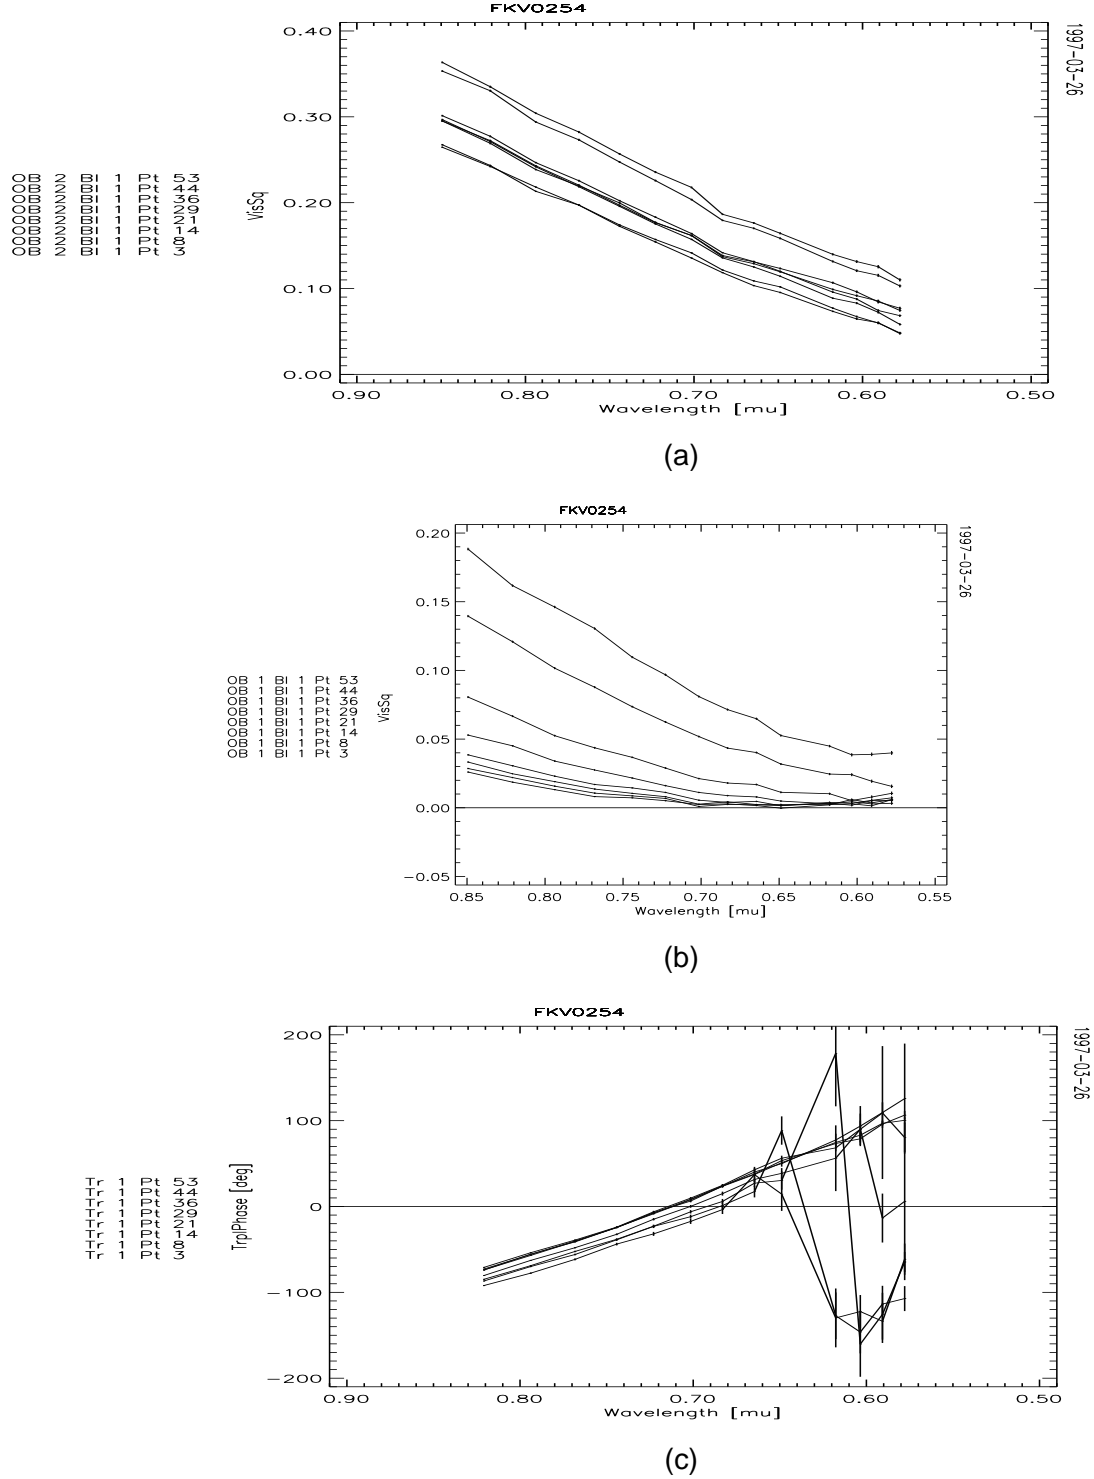


Figure 15.1: Raw visibilities and closure phases of ϵ Gem from the NPOI. (a) Data from the 19 m baseline. (b) Data from the 38 m baseline. (c) Closure phase. Note that the point at $\lambda = 0.62 \mu\text{m}$, $\phi_{\text{cl}} = 170^\circ$ should be shifted down 360° —which is permissible because the phase is determined modulo 360° —to be consistent with the other points.

star. Whether the diameter we fit is a uniform disk or a limb-darkened disk depends only on the model we choose, not on the data, since the shapes of the two models are virtually identical within the first zero of V^2 .

On the other hand, if we pretend that the 37.5 m baseline data shown in Figure 15.1(b) is our interferometer, we can do the observations only at some hour angles. The best of the scans has FASN $\simeq 5$, and still just reaches $V^2 = 0$, demonstrating the utility of wavelength bootstrapping. Andreas Quirrenbach used just this approach on the Mark III interferometer to measure the diameter—and the limb darkening—of Arcturus (Quirrenbach *et al.*, 1996). However, the scan that best detects the zero of V^2 has FASN $\simeq 1.5$, probably too low to fringe track. The fact that we were able to get these data anyway demonstrates the utility of combining wavelength and baseline bootstrapping.

What was B_{track} for these observations? Assuming that we need FASN ≥ 2.5 and using the observed number of photons per frame, we need $V^2 \geq 0.042$ averaged across the frame. The system visibility (i.e., the observed V^2 of a point source, given the imperfections in the instrument and the effects of the atmosphere) was ~ 0.6 , so we need to know the spatial frequency at which $V_*^2 = 0.07$. For a star with this amount of limb darkening, that visibility occurs at $\mathbf{x} = 3.15$. Taking $\lambda = 750$ nm as the representative wavelength, we find $B_{\text{track}} = 32$ m. That is satisfyingly close to the actual baseline length of 37.5 m, since it implies that we could fringe track on that baseline only at hour angles great enough to make the projected baseline less than B_{track} .

Since we actually have a three-element interferometer, we have one more piece of information to help us fit the angular diameter: the closure phase, ϕ_{cl} . Figure 15.1(c) shows the 180° phase jumps at the wavelengths at which $V^2(\lambda)$ reaches zero on the longest baseline. The spatial frequency of that V^2 null is constant; the wavelength at which it occurs changes as the projection of the baseline on the sky changes. In some circumstances, the signal-to-noise ratio for ϕ_{cl} is better than that for V^2 or $V^{(3)}$ (see Section 15.4), so the phase is the best indication of the angular diameter.

15.2.2 Stellar Surface Structure

The hardest problem is imaging a stellar surface, because all the information about surface structure is far out in the (u, v) plane where the fringes are too weak to track. By definition, any surface structure shows up at spatial frequencies that are higher than that of the first zero of V^2 , i.e., for $\mathbf{x} > 3.83$, where \mathbf{x} is calculated using θ_{UD} , the best fit uniform-disk diameter.

So do the data shown in Figure 15.1 tell us anything about surface structure? If we include limb darkening in this category, the answer is that they do tell us something, but not very much. For some scans, the data extend out to somewhere in the neighborhood of the first maximum after the null, but the signal-to-noise ratio of the visibilities is not good enough

there to determine the limb darkening. And certainly the data do not extend to large enough values of \mathbf{x} to detect a spot.

How many array elements do we need to detect a spot? It depends on the size of the spot, of course, so let's say we're trying to detect a spot whose diameter is 20% of the stellar diameter. You can make a rough estimate by imagining that you're trying to measure the diameter of a "star" that is the size of the spot, but with negative flux: you should need about five times the baseline needed for the measurement of the real star, which calls for six array elements.

It is clear that to do these observations, bootstrapping is indispensable. You want to make maximum use of baseline bootstrapping by choosing an array that has the longest possible chain of fringe-tracking baselines. You can reduce the necessary baseline length somewhat by using wavelength bootstrapping. If you're willing to settle for a detection of the spot rather than a measurement of its size, you may be able to reduce the baseline a bit more.

15.2.3 Binary-Star Separation

The choice of array configuration to observe binary stars requires a little more thought. The binary separation $|\rho|$ is larger than either star, but $B_{\text{track}} \simeq \lambda/\theta$; therefore, B_{track} is always more than long enough to resolve the binary, and keeping the array phased up is not the limiting factor. The choice of configuration and/or observing schedule is determined by the pattern of $V^2(u, v)$ values due to the binary. This pattern is formed by the fringe patterns of the two stars beating against one another, causing an oscillation of V^2 whose maxima in the (u, v) plane lie along lines perpendicular to the vector $\vec{\rho}$ separating the stars. The distance between crests in the (u, v) plane is $1/|\rho|$, and one of the crests passes through the origin.

Our goal is to sample this pattern of maxima well enough to determine its spacing and orientation, and thus to determine $\vec{\rho}$. We can sample the pattern radially, by observing in multiple spectral channels on each baseline, or circumferentially, by using Earth rotation to measure V^2 along an elliptical arc in the (u, v) plane at each wavelength, or both. With a three-element array, each observation gives us V^2 along three rays in different directions in the (u, v) plane, which will usually give us a unique solution for $\vec{\rho}$ even without using Earth rotation. Obviously, at least one baseline should be longer than $\lambda/2|\rho|$ or so; otherwise, the binary is unresolved. Not so obviously, at least one baseline should be shorter than $(\lambda^2/\delta\lambda)(|\rho|/n)$, where n is maybe 4 or so; otherwise, the spectral channels (width $\delta\lambda$) will sample too much of the oscillation, washing it out and failing to detect the binary. A more familiar way of looking at this constraint is as a condition on the fractional bandwidth of the channels. You want $\lambda/\delta\lambda \lesssim (\lambda/B)(|\rho|/n)$; otherwise, you get too much bandwidth smearing to see the binary.

The visibility data we end up with are shown well in Figure 14.3: we see an oscillation of V^2 as a function of (u, v) distance. As Dave said, there is a lot of information in these

data. The envelope of V^2 tells us the diameter. The depth of the troughs tells us the magnitude difference between the components, Δm : when $\Delta m = 0$, the troughs go down to zero visibility, while nonzero Δm produces shallower troughs.

We also get ϕ_{cl} data if there are three baselines or more. The detailed shape of $\phi_{\text{cl}}(u, v)$ also contains information about Δm . When Δm is close to zero, ϕ_{cl} in fact is the better indicator of Δm .

15.2.4 Circumstellar Material

Yet another type of observation is imaging extended circumstellar structure. John Monnier showed several examples (Figure 13.13), and Dave Mozurkewich also showed a model of the shape of the visibility with (u, v) distance (Figure 14.4). The information about the extended emission is in the bumps and wiggles in both the visibilities and the closure phases at low spatial frequencies. As with a binary, extended structure is relatively easy to deal with, because the star, which is the phase reference, is relatively compact compared to the structure you want to image, so once again, we may choose a configuration in which all the baselines can be shorter than B_{track} , and in which we have a variety of spacings to optimize the coverage.

15.3 Closure Phase and Calibration

So much for selecting the array. Now I will show what can happen to closure phases with some more data from the NPOI. There are two features that distinguish the NPOI data from the Keck aperture-masking results that John Monnier showed. First, we have only one triangle, since we are currently working with only three array elements, while John was working with 15 or 21 holes (sub-apertures). Second, we take data in 32 spectral channels, so we get 32 closure phases from our three-element array. Typically only the first 16 of those have sufficient signal-to-noise to be useful.

Figure 15.2 shows some uncalibrated NPOI ϕ_{cl} data from unresolved sources. These are stars that are small and single, so we're not getting any data from beyond the first null. Now, closure phase is supposed to be a good observable, that is, it's unaffected by telescope errors (including the atmosphere above the telescopes), and it ought to be zero for an unresolved single source. But ϕ_{cl} runs from -180° to -100° . So what's going on here?

The answer of course is that ϕ_{cl} is immune to telescope-based errors, but not to baseline-based errors. With the type of beam combiner we have at the NPOI [see Figure 14.12(a)], there are opportunities for baseline-based errors to arise. For instance, the light from input beam C is divided into reflected and transmitted beams at the first beam splitter. The reflected beam transits the first beam splitter twice before it combines with beam A at the second beam splitter, while the transmitted beam transits each beam splitter once before it combines with beam B . If the two beam splitters are not exactly the same thickness, the




Figure 15.2: Raw closure phases of unresolved stars observed with the NPOI as a function of time and wavelength. The wavelength scale increases along the axis coming out of the page, running from 550 to 850 nm. The closure phase axis runs from -180° to -80° .




Figure 15.3: Calibrated closure phases as a function of time for the reddest channel of the NPOI, $\lambda = 850$ nm.

two parts of beam C experience different delays. The two baselines for which C is one of the elements will have different phase contributions, creating a baseline-based error.

Another possible source of closure-phase error is nonclosure of apparently closed triangles of array elements. In general, any three apertures A , B , and C will have different wavefront distortions across them. Some of those distortions are atmospheric and eventually average to zero, and some are due to optical imperfections and are constant. The parts of apertures A and B that correlate well—call them subapertures A_1 and B_1 —and that therefore have the dominant contributions to ϕ_{cl} will be different from the parts of B and C that correlate well (B_2 and C_2) and from C and A that correlate well (C_3 and A_3). Now, $A_1\vec{B}_1 + B_2\vec{C}_2 + C_3\vec{A}_3 \neq 0$, i.e., they don't form a triangle, so the phase that we measure on this trio of array elements is not really a closure phase.

Whatever the causes, Figure 15.2 shows that ϕ_{cl} changes smoothly with time and wavelength. In Figure 15.3, we see calibrated ϕ_{cl} from the $\lambda = 850$ nm (i.e. the reddest, and therefore the best behaved) channel. The rms deviation from zero is only about 1.5° , demonstrating how accurate the calibration can be.

In this case, we generated phase corrections by fitting a paraboloid in the wavelength dimension and by smoothing in the time dimension. One of the improvements that we are investigating is generating a correction as a function of wavelength by determining the amount of air-path and glass-path mismatches within the beam combiner needed to produce the observed curvature of phase versus wavelength. Since the beam combiner is in a temperature-stabilized room, that correction should be quite stable during a night, and could significantly improve the phase calibration.

15.4 Model Fitting and Imaging

Once we have calibrated data in hand, the next thing is to interpret it. There are essentially two ways of doing this: model fitting or imaging. In fitting a model, of course, we minimize the difference between the data and the model predictions.

Figure 15.4 shows an example of fitting a model of θ^2 Tau (a spectroscopic binary with $\Delta m = 1.1$ mag) to the calibrated V^2 and ϕ_{cl} data of 23 January 1998, when $|\rho| = 9.0$ mas. A model including the orbital elements and Δm was fitted to 30 nights of Mark III data and six nights of NPOI data. The superimposed model curves in Figure 15.4 are derived from this 30-night data set rather than from this night alone. Conveniently, a single Δm for all wavelengths fits this binary well: it comprises two A stars, one of which has evolved off the main sequence just enough to lie just above the other one on the H-R diagram.

The V^2 data clearly show an oscillation as a function of spatial frequency, with maxima occurring at different spatial frequencies for different scans as the Earth rotates the projected baseline against the pattern of V^2 maxima.

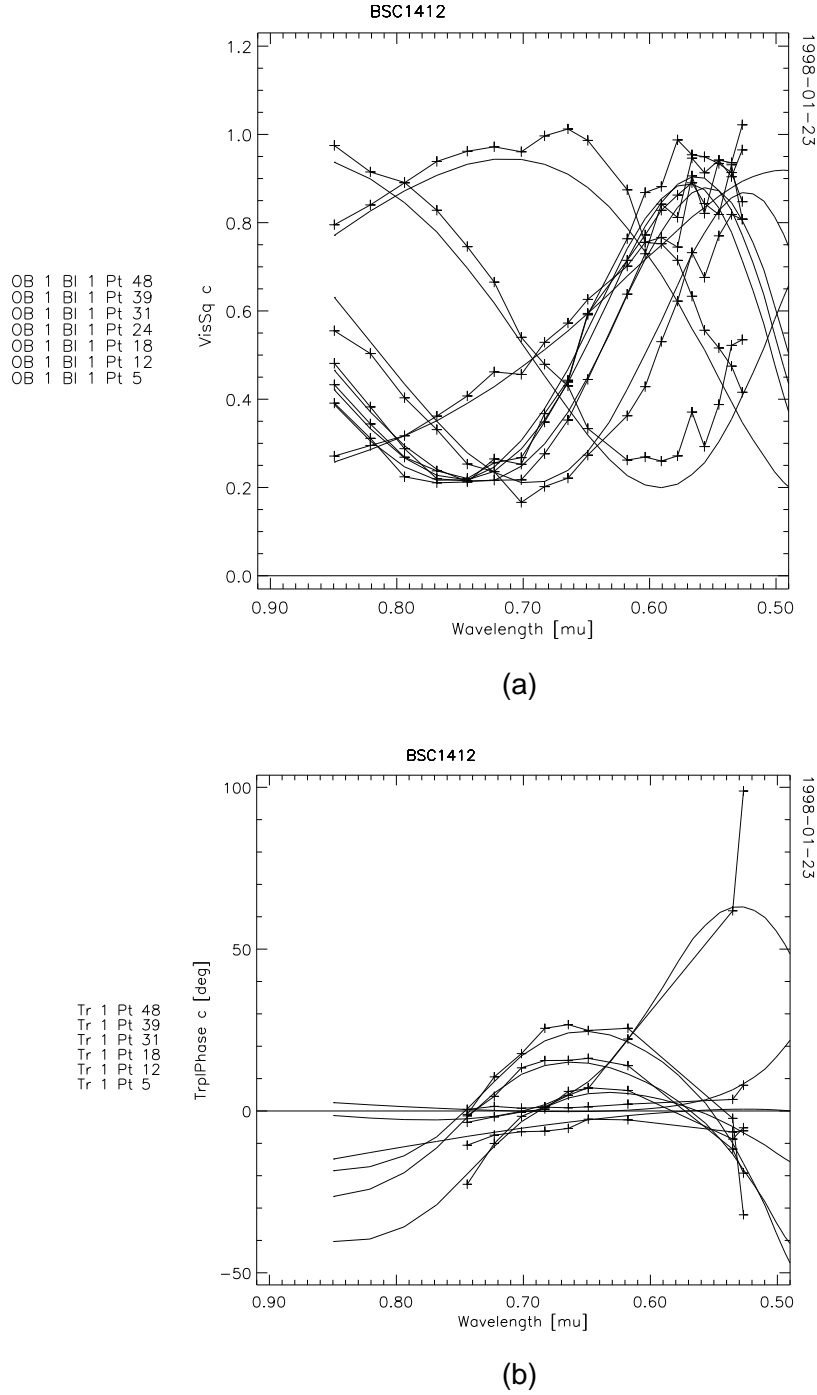


Figure 15.4: NPOI data on θ^2 Tau on 23 January 1998, with model curves superimposed. Error bars have been suppressed for clarity. The data are shown as a function of decreasing wavelength. For V^2 , this corresponds to increasing spatial frequency. (a) $V^2(\lambda)$ data for seven scans on the 37.5 m baseline. (b) $\phi_{cl}(\lambda)$ data for six scans. The four reddest channels are missing due to bad detectors on one of the baselines.

The reason for the shapes of the ϕ_{cl} curves is less intuitively obvious. As we move from short wavelengths to long, the visibilities are sampled along rays in the (u, v) plane. Whenever we cross a minimum of V^2 along one of those rays, the phase on that baseline—and thus ϕ_{cl} —changes by an amount that depends on Δm : for zero magnitude difference, the phase change is 180° . As Δm increases, the size of the phase jump decreases, and in addition, the transition becomes less abrupt.

For an example of modeling the triple product, a different form of the data, I return to the example of ϵ Gem. Figure 15.5 shows the calibrated amplitudes and phases of the triple product as a function of wavelength, with model curves superimposed.


You will recall that the closure phase is the phase of the product of the complex visibilities on three baselines forming a triangle (Section 13.2). Since V data that are measured beyond a null are 180° out of phase with those inside the null, the values of ϕ_{cl} show a step function at the wavelength where one of the contributing baselines crosses the null. The amplitude of the complex product is the triple amplitude $V^{(3)}$. When one of the baselines is sampling near the first null, $V^{(3)}$ may offer a better signal-to-noise ratio than does that V^2 , as appears to be the case in Figure 15.5. In this instance a model diameter for the star should be compared to $V^{(3)}$ rather than to V^2 data.

The other route to interpreting the data is making an image from it. The advantage of imaging is that you're not restricted to your preconceptions about what the source might look like. Rather than impose any constraints on the data, we use the data themselves to tell us what the star looks like.


There are two difficulties with this approach. The first is that we need complex visibilities to Fourier transform into an image, but instead we have V^2 and ϕ_{cl} . We would like to recreate the $N(N-1)/2$ baseline phases, but there are only $(N-1)(N-2)/2$ independent closure phases. John Monnier described in Section 13.3.4 how one starts with an initial guess for the baseline phases and then uses the radio-astronomy based self-calibration and CLEAN routines to iterate to a best fit between the V^2 and ϕ_{cl} data and the model. This technique is still being worked out, and mostly with tools built for radio astronomy that assume that we're working with complex visibilities. There is a good deal of development to be done in this direction.

The second difficulty is that we must deconvolve the point-spread function of the array from the image. This is an art that has been extensively developed in radio astronomy, as well as in a number of other fields. It is here that the distinction between modeling and imaging starts to break down. For instance, the CLEAN algorithm treats an image as an ensemble of delta functions. The model is convolved with the point-spread function to produce the final product. People think of the process of producing a CLEANed, self-calibrated result as imaging, but it is imaging with a strong modeling component.

At NPOI we have produced images from data of V^2 and ϕ_{cl} by first assigning ϕ_{cl} to one of the baselines, then self-calibrating and CLEANing. Because we have been working with



(a)



(b)

Figure 15.5: Calibrated closure phases and triple amplitudes of ϵ Gem compared to a limb-darkened model with $T_{\text{eff}} = 4730$ K and $\log g = 1.5$ (Evans & Teays, 1996), and $\theta_{\text{LD}} = 4.89$ mas (Pauls *et al.*, 1998). (a) Closure phases and model curves. (b) Triple amplitudes and model curves.

only three array elements, we don't have the problem that John Monnier had in his Keck aperture masking data, of choosing which set of independent closure phases to use. But with only three array elements commissioned so far, we are also limited to much simpler images, those of binary stars. It is encouraging, however, that even in these simple images, we have attained a dynamic range of between 100:1 and 200:1, which is about the same as in the aperture masking images.

15.5 Current Capabilities and Limitations

The current capabilities and limitations of optical interferometry include high angular resolution, low sensitivity, moderate complexity, and moderate dynamic range. The highest resolutions that have been attained by long-baseline interferometers are between 1 and 2 mas for measurements of stellar angular diameters, and 2 mas or a little more for separations of binary stars, far outstripping any other techniques available at visual and infrared wavelengths. But interferometers demand a lot of light. The magnitude limit at visual wavelengths is currently about 6^m for long-baseline interferometers, and about 3^m for aperture masks.

The current long-baseline optical interferometers can produce images of relatively low complexity because they sample a small number of (u, v) points; even with the NPOI's 32 spectral channels on each of three baselines, we get only 33% of the phase information. For comparison, the VLA has 351 baselines. The Keck aperture masks in the observations John Monnier described used 210 baselines and yield 190 independent closure phases, or 90% of the phase information available before the light hit the atmosphere. The increased information content is clear in the complexity of the resulting images.

The dynamic range in images from both long-baseline and aperture-masking interferometry is in the range of 100:1 to 200:1, largely determined by the precision of the calibration. Improving the calibration is difficult, but a promising approach is to use single-mode fibers as spatial filters, as is being tried by the FLUOR group using IOTA and by the group at PTI (cf. Traub, Section 3.4.2).

15.6 Future Developments

There are two directions in which the field of optical interferometry is growing. One is toward aperture masking of the kind that the Cambridge group and the Keck group are doing, using masks with a large number of holes and thus taking data from a large number of baselines and retrieving upwards of 80% of the phase data. There is a lot of interesting work to be done with this technique, even though it has two distinct limitations. One is that the maximum baseline is 10 m; the other is that the images are monochromatic.

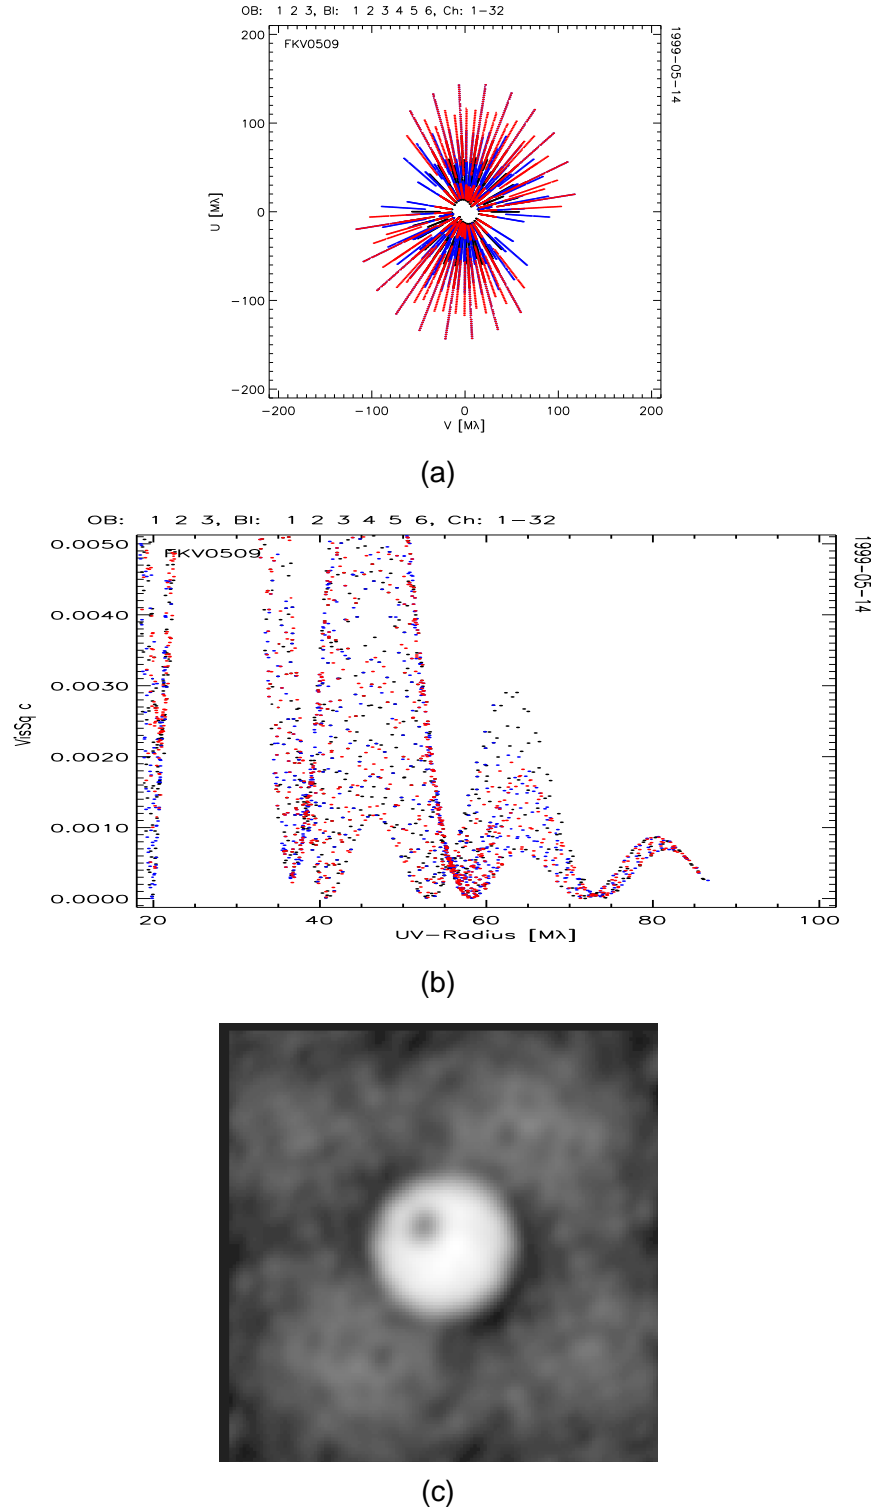


Figure 15.6: Simulation of a 12-mas limb-darkened star with a spot imaged with a six-element array in 32 spectral channels. (a) (u, v) coverage. (b) V^2 vs. $\sqrt{u^2 + v^2}$ beyond the first null, before noise addition. (c) Image restored with a maximum-entropy routine. [Simulation courtesy of Christian Hummel.]

The other direction is toward increasing the number of elements and the maximum baseline in such long-baseline interferometers as COAST, the NPOI, the CHARA array, or the GI2T. Currently, COAST has five elements working and can use four at any given time; the NPOI has three, but should have six by the end of 2000 with a maximum baseline of 65 m (and ultimately 437 m); CHARA has had first light on one baseline, and will ultimately have as many as eight elements and a 400 m maximum baseline; and IOTA and GI2T will have three elements within about the same period, with somewhat shorter baselines. The longest baseline may eventually be that of SUSI, at 640 m, although for the next few years it is likely to remain a two-element interferometer.

These long-baseline interferometers overcome both of the limitations of the aperture masking technique. First, baseline length is no longer limited by the size of a single telescope. Second, separating the light-gathering from the imaging functions of the instrument makes it possible to take data in a large number of spectral channels. Increasing the number of data channels can offset the advantage the aperture masking technique has in number of baselines.

Figure 15.6 shows an indication of the image complexity that a six-element multi-channel array will be capable of. We simulated observations of a 12-mas limb-darkened star with 5% noise added to the V^2 data. We used six elements in one of the baseline-bootstrapped arrays that Dave Mozurkewich showed (Figure 14.8), with three telescopes evenly spaced along each of two arms of the NPOI Y-shaped imaging array, and simulated 12 scans, using Earth rotation to improve the (u, v) coverage. We then deconvolved the image with a maximum-entropy routine. The result is comparable to the images from aperture masking observations.

Two and three element optical interferometers have already shown their ability to measure stellar diameters and binary orbits. The multiple-baseline interferometers that are coming into operation over the next few years should allow us to see the details of what is happening on and around the surfaces of stars.

References

- N.R. Evans and T.J. Teays, “Temperatures of δ Cep and nonvariable supergiants,” *Astrophys. J.* **112**, 761–771 (1996).
- T.A. Pauls, D. Mozurkewich, J.T. Armstrong, C.A. Hummel, J.A. Benson, and A.R. Hajian, “Observations of stellar limb darkening with the Navy Prototype Optical Interferometer,” in *Astronomical Interferometry*, R.D. Reasenberg, ed. Proc. SPIE, **3350**, 467–470 (1998).
- A. Quirrenbach, D. Mozurkewich, D.F. Buscher, C.A. Hummel, and J.T. Armstrong, “Angular diameter and limb darkening of Arcturus,” *Astron. Astrophys.* **312**, 160–166 (1996).

